

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 042390.P8918 Total Pages 2

First Named Inventor or Application Identifier Salvador Palanca

Express Mail Label No. EL546447249US

ADDRESS TO: Assistant Commissioner for Patents
Box Patent Application
Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. X Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)
2. X Specification (Total Pages 28)
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claims
 - Abstract of the Disclosure
3. X Drawings(s) (35 USC 113) (Total Sheets 8)
4. x Oath or Declaration (Total Pages 6) (unsigned)
 - a. Newly Executed (Original or Copy)
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed) (**Note Box 5 below**)
 - i. DELETIONS OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (useable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Microfiche Computer Program (Appendix)

7. ☐ Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
a. ☐ Computer Readable Copy
b. ☐ Paper Copy (identical to computer copy)
c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & documents(s))
9. ☐ a. 37 CFR 3.73(b) Statement (where there is an assignee)
☒ b. Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ a. Information Disclosure Statement (IDS)/PTO-1449
☐ b. Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
14. ☐ a. Small Entity Statement(s)
☐ b. Statement filed in prior application, Status still proper and desired
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. ☒ Other: Copy of postcard with Express mail stamp

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP)
of prior application No: ☐

18. Correspondence Address

☐ Customer Number or Bar Code Label _____
(Insert Customer No. or Attach Bar Code Label here)
or

☒ Correspondence Address Below

NAME Sanjeet Dutta (46,145) Sanjeet Dutta 6/30/2000
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard
Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EL546447249US

Date of Deposit: _____

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Trademarks, 2900 Crystal Drive, Arlington, Virginia 22202-3513.

Monique M. Pearson

(Typed or printed name of person mailing paper or fee)

[Signature]

(Signature of person mailing paper or fee)

6/30/00

(Date signed)

Serial/Patent No.: Unassigned

Filing/Issue Date: June 30, 2000

Client: Intel Corporation

Title: METHOD AND APPARATUS FOR CACHE REPLACEMENT FOR A MULTIPLE VARIABLE-WAY ASSOCIATIVE CACHE

BSTZ File No.: 042390.P8918

Atty/Secty Initials: MJM/SKD/mmp

Date Mailed: June 30, 2000

Docket Due Date: ***

The following has been received in the U.S. Patent & Trademark Office on the date stamped hereon:

- | | | |
|--|--|--|
| <input type="checkbox"/> Amendment/Response (____ pgs.) | <input checked="" type="checkbox"/> Express Mail No.: <u>EL546447249US</u> | <input checked="" type="checkbox"/> Check No. <u>36185</u> |
| <input type="checkbox"/> Appeal Brief (____ pgs.) (in triplicate) | <input type="checkbox"/> _____ Month(s) Extension of Time | Amt: <u>864.00</u> |
| <input checked="" type="checkbox"/> Application - Utility (<u>29</u> pgs., with cover and abstract) | <input type="checkbox"/> Information Disclosure Statement & PTO 144 (____ pgs.) | <input type="checkbox"/> Check No. _____ |
| <input type="checkbox"/> Application - Rule 1.53(b) Continuation (____ pgs.) | <input type="checkbox"/> Issue Fee Transmittal | Amt: _____ |
| <input type="checkbox"/> Application - Rule 1.53(b) Divisional (____ pgs.) | <input type="checkbox"/> Notice of Appeal | |
| <input type="checkbox"/> Application - Rule 1.53(b) CIP (____ pgs.) | <input type="checkbox"/> Petition for Extension of Time | |
| <input type="checkbox"/> Application - Rule 1.53(d) CPA Transmittal (____ pgs.) | <input type="checkbox"/> Petition for _____ | |
| <input type="checkbox"/> Application - Design (____ pgs.) | <input checked="" type="checkbox"/> Postcard | |
| <input type="checkbox"/> Application - PCT (____ pgs.) | <input type="checkbox"/> Power of Attorney (____ pgs.) | |
| <input type="checkbox"/> Application - Provisional (____ pgs.) | <input type="checkbox"/> Preliminary Amendment (____ pgs.) | |
| <input type="checkbox"/> Assignment and Cover Sheet | <input type="checkbox"/> Reply Brief (____ pgs.) | |
| <input checked="" type="checkbox"/> Certificate of Mailing <u>Express Mail</u> | <input type="checkbox"/> Response to Notice of Missing Parts | |
| <input checked="" type="checkbox"/> Declaration & POA (<u>6</u> pgs.) (unsigned) | <input type="checkbox"/> Small Entity Declaration for Indep. Inventor/Small Business | |
| <input type="checkbox"/> Disclosure Docs & Orig & Copy of Inventor's Signed Letter (____ pgs.) | <input checked="" type="checkbox"/> Transmittal Letter, in duplicate | |
| <input checked="" type="checkbox"/> Drawings: <u>8</u> # of sheets includes <u>11</u> figures | <input checked="" type="checkbox"/> Fee Transmittal, in duplicate | |

☒ Other: Copy of postcard with Express Mail stamp

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**Method and Apparatus for Cache Replacement For
A Multiple Variable-Way Associative Cache**

Inventors: Salvador Palanca
Subramaniam Maiyuran

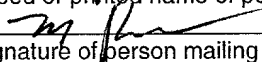
Prepared by:
Blakely, Sokoloff, Taylor & Zafman
1279 Oakmead Parkway
Sunnyvale, California 94086
(408) 720-8300

"Express Mail" mailing label number EL546447249US

Date of Deposit 6/30/00

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Monique M Pearson
(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

Method and Apparatus for Cache Replacement in A Multiple Variable-Way Associative Cache

FIELD OF THE INVENTION

5 The present invention relates to the field of computer systems. In particular, the present invention relates to a method and apparatus for cache replacement in a multiple variable-way associative cache.

BACKGROUND OF THE INVENTION

10 Caches are commonly used to temporarily store values that might be repeatedly accessed by a processor, in order to speed up processing by avoiding the longer step of loading the values from main memory such as random access memory (RAM).

15 A cache has many "blocks" which individually store the various instructions and data values. The blocks in any cache are divided into groups of blocks called "sets." A set is the collection of cache blocks that a given memory block can reside in. For any given memory block, there is a unique set in the cache that the block can be mapped into, according to preset mapping functions. The number of blocks in a set is referred to as the associativity of the cache,
20 e.g., 2-way set associative means that, for any given memory block there are two blocks in the cache that the memory block, can be mapped into; however, several different blocks in main memory can be mapped to any given set. A 1-way set associative cache is direct mapped; that is, there is only one cache block

that can contain a particular memory block. A cache is said to be fully associative if a memory block can occupy any cache block, i.e., there is one set, and the address tag is the full address of the memory block.

An exemplary cache line (block) includes an address-tag field, a
5 state-bit field, an inclusivity-bit field, and a value field for storing the actual instruction or data. The state-bit field and inclusivity-bit field are used to maintain cache coherency in a multiprocessor computer system. The address tag is a subset of the full address of the corresponding memory block. A compare match of an incoming effective address with one of the tags within the address-tag field
10 indicates a cache "hit." The collection of all of the address tags in a cache (and sometimes the state-bit and inclusivity-bit fields) is referred to as a directory, and the collection of all of the value fields is the cache entry array.

When all of the blocks in a set for a given cache are full and that cache receives a request, with a different tag address, whether a "read" or
15 "write," to a memory location that maps into the full set, the cache must "evict" one of the blocks currently in the set. The cache chooses a block by one of a number of means known to those skilled in the art (least recently used (LRU), random, pseudo-LRU, etc.) to be evicted. If the data in the chosen block is modified, that data is written to the next lowest level in the memory hierarchy
20 which may be another cache (in the case of the L1 or on-board cache) or main memory (in the case of an L2 cache, as depicted in the two-level architecture of FIG. 1). By the principle of inclusion, the lower level of the hierarchy will already

have a block available to hold the written modified data. However, if the data in the chosen block is not modified, the block is simply abandoned and not written to the next lowest level in the hierarchy. This process of removing a block from one level of the hierarchy is known as an "eviction." At the end of this process,

5 the cache no longer holds a copy of the evicted block.

This ratio of available blocks for instruction versus data is not, however, always the most efficient usage of the cache for a particular procedure. Many software applications will perform better when run on a system with split I/D caching, while others perform better when run on a flat, unified cache (given

10 the same total cache space). In the instances where the cache I/D ratio is not particularly close to the actual ratio of instruction and data cache operations, there are again a troubling number of evictions.

A cache replacement algorithm determines which cache block in a given set will be evicted. For example, an 8-way associative cache might use an

15 LRU unit which examines a 7-bit field associated with the set.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are included as part of the present specification, illustrate the presently preferred embodiment of the present invention and together with the general description given above and the detailed description of the preferred embodiment given below serve to explain and teach the principles of the present invention.

Figure 1 illustrates an integrated multi- processor computer system;

Figure 2 illustrates a cache memory having a sharing mode and non sharing mode;

Figure 3 illustrates a pseudo-LRU algorithm for an 8-way set associative cache;

Figure 4A illustrates cache device 400 in non-sharing mode at time t_0 ;

Figure 4B illustrates cache device 400 in sharing mode at time t_1 ;

Figure 4C illustrates cache device 400 in sharing mode at time t_2 ;

Figure 4D illustrates cache device 400 in non-sharing mode at time t_3 ;

Figure 5 illustrates a multiple pseudo-LRU replacement mechanism for an 8-way cache way subdivided into 6-way and 2-way set associativities;

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

A method and apparatus for cache replacement in a multiple variable-way associative cache. The method according to the present techniques partitions a cache array dynamically based upon requests for memory from an
5 integrated device having a plurality of processors.

In the following description, for purposes of explanation, specific nomenclature is set forth to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention. For example,
10 the present invention has been described with reference to documentary data. However, the same techniques can easily be applied to other types of data such as voice and video.

Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on
15 data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical
20 manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven

convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar
5 terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to
10 the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or
15 display devices.

The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a
20 computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access

memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

Figure 1 illustrates an integrated multi- processor computer system. System 100 may have one or more processing units, such as Central Processing unit (CPU) 111 and graphics processor 113. CPU 111 and graphics processor 113 are integrated with memory controller 112 into integrated multi-processing device 110. Although described as fully integrated, device 110 could be broken into individual components in alternate embodiments.

Included in device 111 is level one cache 120 which is implemented using high speed memory devices. L1 cache 120 is a small on-board cache. In one embodiment the L1 cache 120 may be only 64 kilobytes. Connected to CPU 111 and graphics processor 113 is a level 2 cache 130. In

one embodiment L2 cache 130 is considerably larger than L1 cache 120, and may be 512 kilobytes. L2 cache 130 supports L1 cache 120. Although **Figure 1** depicts only a two-level cache hierarchy, multi-level cache hierarchies can be provided where there are many levels of interconnected caches.

5 Multi-processor device 110 is connected to bus 170. Also connected to bus 170 are various peripheral devices, such as, input/output (I/O) devices 150 (i.e. a display monitor, keyboard, or permanent storage device), main memory devices 160 (i.e. random access memory (RAM), or firmware 140 (i.e. read only memory (ROM)). Firmware 140 is used to load operating systems,
10 commands, and drivers for I/O devices 150. Memory devices 160 are used by the processors in device 110 to carry out program instructions. Memory controller 112 manages the transfer of data between the processor core and the cache memories, 120 and 130.

 L2 cache 130 acts as an intermediary between main memory 160
15 and L1 cache 120, and has greater storage ability than L1 cache 120, but may have a slower access speed. Loading of data from main memory 160 into multi-processor device 110 goes through L2 cache 130. L2 cache 130 can be subdivided so that processors within device 110 may share L2 cache 130 resources, thus, allowing higher system performance for the same available
20 memory bandwidth. For example, graphics processor 113 and CPU 111 may access L2 cache 130 simultaneously without degrading the bandwidth or latency of CPU 111.

L2 cache 130 operates in two modes. In “non-sharing” mode, L2 cache 130 performs normally and dedicates all its resources to CPU 111. In “sharing” mode, L2 cache 130 dynamically partitions its resource based upon requests from multi-processor device 110. When in sharing mode, CPU 111 perceives one portion of the L2 cache 130, and graphics processor 113 perceives the remainder of the L2 cache 130. For example, in sharing mode when three dimensional graphics applications are run on multi-processor device 110, fifty percent of L2 cache 130 is allocated to CPU 111 and fifty percent of L2 cache 130 is allocated to graphics processor 113. Thus, the cache size allocated for graphics can potentially marginally degrade the performance of CPU 111.

Figure 2 illustrates a cache memory having a sharing mode and non sharing mode. Cache memory 300 is an 4-way set associative cache, having ways 0 - 3 and no sets (A-P) 309-324. In non-sharing mode all sets (A-P) 309-324 are allocated to CPU 111. In sharing mode, sets A-H 309-316 may be allocated to CPU 111, and sets I – P 317-324 may be allocated to graphics processor 113, in one embodiment. In another embodiment, each set could be partitioned. Thus, in sets A – P 309-324 ways 0 - 3 are divided. Thus ways 0_A , and 1_A in set A 309 may be allocated to CPU 111, and ways 2_A and 3_A in set A 309 may be allocated to graphics processor 113. Similarly, sets B – P 310-324 would be divided.

Color, Z, and texture are examples of graphics request supported during cache sharing. In one embodiment, color and Z achieve improved system performance when used with direct-mapped caches. Texture improves system performance when used with multiple-way set associative caches.

5 Thus, way-subdivision and set-subdivision have just been described. Way subdivided means that for a cache of X sets and Y ways, X sets and Y minus U ways are used to store CPU 111 data, while X sets and U ways store graphics processor 113 data, when in sharing mode. Set subdivided means X minus V sets and Y ways are allocated for CPU 111 data, while V sets and Y ways are allocated for graphic processor 113 data, when in sharing mode. In one embodiment, X, V, Y, and U are numbers that are multiples of two.

10 In a way subdivided cache, the number of sets allocated for each request type (i.e. CPU, texture, Z or color) remains constant, while the number of ways decreases. For example, a cache with X sets and Y ways in sharing mode supports simultaneous requests of various types where half of the cache is allocated to CPU 111 transactions, one quarter of the cache 300 is allocated for texture transactions and one quarter of the cache 300 is allocated for color/Z. Consequently, X sets and Y/2 ways are allocated to CPU 111 transactions, X sets and Y/4 ways are allocated for texture transactions and X sets and Y/4 ways are allocated for color/Z transactions.

20 For example, **Figures 3A – 3D** illustrate a way-subdivided cache in two different sharing modes, and non-sharing mode. **Figure 4A** illustrates cache

device 400 in non-sharing mode at time t_0 . Cache device 400 is an 8-way set associative cache with 2 sets (A,B) 401 and 409 and 8 ways $0_A - 7_A$. All ways $0_A - 7_A$ are allocated to CPU 111. **Figure 4B** illustrates cache device 400 in sharing mode at time t_1 . Cache device 400 is way subdivided and configured for three simultaneous accesses. Thus, set A 411 is a 4 way set associative cache having ways $0_A - 3_A$. Set B 412 is a 2 way set associative cache having ways 0_B and 1_B . Set C 413 is a 2 way set associative cache having ways 0_C and 1_C . In one embodiment, Set A 411 is allocated for CPU 111 transactions, set B 412 is allocated for texture transactions, and set C 413 is allocated for color/Z transactions. The remainder 419 of the cache array 400 is way subdivided in a similar fashion.

Figure 4C illustrates cache device 400 in sharing mode at time t_2 .

Cache device 400 is way subdivided and configured for two simultaneous accesses. Thus, set A 421 is a 6-way set associative cache having ways $0_A - 5_A$. Set B 422 is a 2 way set associative cache having ways 0_B and 1_B . In one embodiment, set A 421 is allocated for CPU 111 transactions, and set B 422 is allocated for texture transactions. No ways are allocated for color/Z transactions. In another embodiment, Set A 421 is allocated for CPU 111 transactions, and Set B 422 is allocated for color/Z transactions. No ways are allocated for texture transactions. The remainder 429 of cache array 400 is way subdivided in a similar fashion.

Figure 4D illustrates cache device 400 in non-sharing mode at time t_3 . Cache device 400 has returned to the non-sharing mode of **Figure 4A**, as an 8-way set associative cache with set A 431 and B 439. All ways $O_A - 7_A$ are allocated to CPU 111 as well.

5 Although described with respect to way subdivision, cache device 400 could be set subdivided in alternate embodiments.

Table 1 shows different configurations and the resulting number of ways and sets in sharing mode for a cache with X sets and Y ways. It is assumed three simultaneous transactions; therefore, each configuration has
10 three numbers representing the cache size allocated to each request type. For example, $\frac{1}{2} - \frac{1}{4} - \frac{1}{4}$ means that in sharing mode, $\frac{1}{2}$ of the cache is used to store CPU data, $\frac{1}{4}$ for texture data and the last $\frac{1}{4}$ for color/Z data. When not in sharing mode, the whole cache is allocated to CPU transactions.

| | CONFIGURATION | | | |
|-----------------|---|---|---|---|
| | 1 - 0 - 0 (not sharing mode) | $\frac{1}{2} - \frac{1}{4} - \frac{1}{4}$ (sharing mode) | $\frac{3}{4} - \frac{1}{4} - 0$ (sharing mode) | $\frac{3}{4} - 0 - \frac{1}{4}$ (sharing mode) |
| Way subdivision | CPU: X sets, Y ways Texture: 0 Color/Z: 0 | CPU: X sets, Y/2 ways Texture: X sets, Y/4 ways Color/Z: X sets, Y/4 ways | CPU: X sets, 3Y/4 ways Texture: X sets, Y/4 ways Color/Z: 0 | CPU: X sets, 3Y/4 ways Texture: 0 Color/Z: X sets, Y/4 ways |
| Set subdivision | CPU: X sets, Y ways Texture: 0 Color/Z: 0 | CPU: X/2 sets, Y ways Texture: X/4 sets, Y ways Color/Z: X/4 sets, Y ways | CPU: 3X/4 sets, Y ways Texture: X/4 sets, Y ways Color/Z: 0 | CPU: 3X/4 sets, Y ways Texture: 0 Color/Z: X/4 sets, Y ways |

15 Table 1: number of sets and ways allocated to each request type for different configurations

As shown in Table 1, when using way subdivision, CPU way associativity decreases in sharing mode; i.e., number of ways in the cache portion allocated to CPU transaction in sharing mode is less than way associativity in non-sharing mode. On the other hand, using set subdivision, way associativity can be maintained constant in sharing mode with no LRU array growth and minimal die size impact.

Table 2 shows the resulting cache sizes and pseudo-LRU algorithms after switching to sharing mode for three different configurations. Three simultaneous accesses; i.e., CPU, texture and color/Z, are assumed in sharing mode.

| | | CONFIGURATION | | | |
|-----------------|---------|---------------------------------|--|-------------------------------------|-------------------------------------|
| | | 1 – 0 – 0 (not sharing mode) | 1/2 – 1/4 – 1/4 (sharing mode) | 3/4 – 0 – 1/4 (sharing mode) | 3/4 – 1/8 – 1/8 (sharing mode) |
| Set subdivision | CPU | X sets, 8 ways | X/2 sets, 8 ways | 3X/4 sets, 8 ways | 3X/4 sets, 8 ways |
| | Texture | 0 | X/4 sets, 8 ways | 0 | X/8 sets, 8 ways |
| | Color/Z | 0 | X/4 sets, 8 ways | X/4 sets, 8 ways | X/8 sets, 8 ways |
| Way subdivision | CPU | X sets, 8 ways | X sets, 4 ways LO=0 L2, L5 and L6 unused | X sets, 6 ways L2=0 L6 unused | X sets, 6 ways L2=0 L6 unused |
| | Texture | 0 | X sets, 2 ways Use L5 | 0 | X sets, 1 way No LRU |
| | Color/Z | 0 | X sets, 2 ways Use L6 | X sets, 2 ways Use L6 | X sets, 1 way No LRU |

Table 2

As shown in Table 2, when using set subdivision in sharing mode, the number of ways allocated to each request type is constant (i.e., it is always 8

ways, regardless of whether the cache is in sharing mode). Therefore, the same LRU algorithm as when not in sharing mode can be used. When using way subdivision, the number of sets remains constant. Consequently, in sharing mode, a single LRU array has to support several pseudo-LRU algorithms.

- 5 **Figure 3** illustrates a pseudo-LRU algorithm for an 8-way set associative cache. Given an 8-way set associative cache as described in Table 2, which uses a pseudo-LRU replacement algorithm, such as that shown in **Figure 3**, when switching to sharing mode using way subdivision, in configuration $\frac{1}{2} - \frac{1}{4} - \frac{1}{4}$, the LRU bit L0 is hard-coded to 0 for every cache set.
- 10 CPU requests have now a 4-way set associative cache, and they use similar LRU algorithms as described herein, but only with LRU bits L1, L3 and L4. Texture and color/Z requests are stored in a 2-way set associative cache each. They use LRU bits L5 and L6, respectively. LRU bit L2 is unused in sharing mode. For a 2-way set associative cache, hits to way 0 set the LRU bit value to
- 15 1, and hits to way 1 clear the LRU bit value to 0.

Similarly, the LRU bit L2 is hardcoded to 0 for every cache set, when switching to sharing mode using way subdivision in configuration $\frac{3}{4} - 0 - \frac{1}{4}$. CPU requests use LRU bits L0, L1, L3, L4, and L5 in a 6-way set associative cache. LRU bit L6 is used for a 2-way set associative color/Z cache. Texture

20 requests are not cached. There is no change as far as CPU requests is concerned for configuration $\frac{3}{4} - \frac{1}{8} - \frac{1}{8}$. In the latter configuration, texture and

color/Z are direct-mapped; therefore, no LRU is needed, and LRU bit L6 is unused in sharing mode.

By making the corresponding logic and circuit changes, not only multiple requests, but also multiple configurations can be supported with the existing LRU cache array used to support CPU cache accesses when not in sharing mode. Other configurations aside from those shown in Table 2 can be similarly implemented. Multiple LRU algorithms can also be supported using the same technique described in **Figure 2** for higher associativity caches; i.e., starting from a non-sharing, mode 16-way set associative cache with 15 LRU bits per cache set, 32-way set associative cache with 31 LRU bits per cache set...

The technique used in table 2 to support multiple LRU algorithms on a way subdivided shared cache works best when the ways allocated to each request type in sharing mode is a power of two. In a way subdivided cache, the LRU algorithm for CPU requests in sharing mode in the configuration $\frac{3}{4}$ -0- $\frac{1}{4}$ can be improved for better performance. LRU bit L2 remains hardcoded for every cache set, when in sharing mode. The combination LOL1=11 is illegal in sharing mode. LRU bit L6 is still used for a 2-way set associative cache for color/Z request, when in sharing mode.

In one embodiment, cache device 400 uses the following multiple pseudo-LRU update mechanism for an 8-way set associative cache with seven LRU bits per cache set. The update mechanism indicates which way of a given set will be replaced upon a cache miss to that particular set. A cache miss

occurs when data sought by CPU 111 or graphics processor 113 is not already in the L2 cache, 130 or L1 cache 120 but instead needs to be extracted from main memory 160. Once extracted from main memory 160, the update mechanism determines where to place the data within cache 130 according to an 8-bit code indicating the least recently used (LRU) ways.

For example, suppose cache 400 is operating in sharing mode with 3/4 of cache 400 for CPU transactions, and 1/4 of cache 400 for color/Z transactions with no caching of texture data. Thus, the 8-way set associative cache is transformed into a 6-way set associative cache for CPU 111 transactions and a 2-way set associative cache for color/Z transactions from graphics processor 113.

| New LRU Values | | | | | | | | |
|----------------|----|----|----|----|----|----|----|--------------|
| Entry Hit | L0 | L1 | L2 | L3 | L4 | L5 | L6 | Replaced Way |
| | 0 | 0 | 0 | 0 | 0 | 0 | 0 | = 0 |
| update 0 | 0 | 1 | | 1 | | | | |
| | 0 | 1 | 0 | 1 | 0 | 0 | 0 | = 2 |
| update 2 | 1 | 0 | | | 1 | | | |
| | 1 | 0 | 0 | 1 | 1 | 0 | 0 | = 4 |
| update 4 | 0 | 0 | | | | 1 | | |
| | 0 | 0 | 0 | 1 | 1 | 1 | 0 | = 1 |
| update 1 | 0 | 1 | | 1 | | | | |
| | 0 | 1 | 0 | 1 | 1 | 1 | 0 | = 3 |
| update 3 | 1 | 0 | | | 0 | | | |
| | 1 | 0 | 0 | 1 | 0 | 1 | 0 | = 5 |
| update | 0 | 0 | | | | 0 | | |

| | | | | | | | | |
|-------------|---|---|---|---|---|---|---|-----|
| 5 | | | | | | | | |
| | 0 | 0 | 0 | 1 | 0 | 0 | 0 | = 1 |
| update 1 | 0 | 1 | | 1 | | | | |
| | 0 | 1 | 0 | 1 | 0 | 0 | 0 | = 2 |
| update 2 | 1 | 0 | | | 1 | | | |
| | 1 | 0 | 0 | 1 | 1 | 0 | 0 | = 4 |
| update 4 | 0 | 0 | | | | 1 | | |
| | 0 | 0 | 0 | 1 | 1 | 1 | 0 | = 0 |
| update 0 | 0 | 1 | | 1 | | | | |
| | 0 | 1 | 0 | 1 | 1 | 1 | 0 | = 3 |
| update 3 | 1 | 0 | | | 0 | | | |
| | 1 | 0 | 0 | 1 | 0 | 1 | 0 | = 5 |
| update 5 | 0 | 0 | | | | 0 | | |
| | 0 | 0 | 0 | 1 | 0 | 0 | 0 | = 1 |
| update 1 | 0 | 1 | | 1 | | | | |
| | 0 | 1 | 0 | 1 | 0 | 0 | 0 | = 2 |

Table 3

Figure 5 illustrates a multiple pseudo-LRU replacement

mechanism, corresponding to the cache sharing mode described in the previous

5 paragraph. LRU bit L2 is never used. LRU bit L6 is only used for color/Z

transactions. On CPU transactions, the 7-bit LRU algorithm is reduced to a 5-bit

LRU algorithm (LRU bits L0, L1, L3, L4, and L5). Table 3 shows how CPU

transactions update the LRU entries. Line 1 of table 3 shows that all LRU bits

begin having the value 0. Looking at **Figure 5**, one sees that LRU code

10 00000000 leads to a hit on way 0. Once way 0 is hit, the LRU code is updated

as shown in line 2 of table 3, that shows LRU bit 0 (L0) becomes 0, L1 becomes 1 and L3 become 1. The resulting 7-bit LRU code is 0101000. Referring to **Figure 5**, one can see that the LRU replacement mechanism directs cache 400 to store the data in way 2. The LRU bits are updated as shown in the table 3.

5 The remaining ways are replaced according to table 3 and **Figure 5** as described above.

Cache 400 may also be dynamically converted into a direct – mapped cache. Some graphics transactions, such as color/Z, achieve improved performance using a direct mapped cache. Regardless of whether way or set subdivision is used, and based on the cache size allocated to a particular transaction, the initial partitioning may not yield a direct-mapped cache for the particular transaction, and further conversion may be required to go to direct-mapped when in sharing mode.

10 An exemplary cache line relating the physical address of a memory cell in cache 400 consists of a tag address, a set address, and a byte offset. In one embodiment, an n-way set associative cache is direct mapped by (for N equals 2 to the power M), by expanding the set address by M bits and decreasing the tag portion of the address by M bits. The M least significant bits of the tag become the most significant bits of the set address.

20 **Figure 6** illustrates transforming N-way set associative caches to a direct mapped cache. A 1-way set associative cache is direct mapped; that is, there is only one cache block that can contain a particular memory block. Cache

device 500 is a 2-way set associative cache shown with 2 sets. Set 0 510 and set 1 520 are transformed into direct mapped caches with only one set. The number of sets doubles. LRU bits are meaningless since there is only one possible way to be mapped. Thus, way 0 of set 0 510 becomes set 0 511 having a single way 0. Similarly, way 1 of set 0 510 becomes set 2 512 having a single way 0, as well. Set 1 520 is unwrapped the same way as set 0 510.

Figure 7 illustrates a flow diagram of converting an N-way set

associative cache into a direct mapped cache. The process begins in block 600.

In processing block 610, an N-way set associative cache array where $N=(2 \text{ to the power } M)$ is selected to be converted. The cache array has L sets, where $L=(2 \text{ to the power } K)$. The cache line size is H bytes, where $H=(2 \text{ to the power } J)$.

The cache is byte addressable, with a physical address (PA) of Z bits.

In processing block 620, the physical address for a single access cache is defined to have three components, a tag address, a set address, and a byte offset. Bits Z to J+K are the tag address, bits J+K-1 to J are the set address, and bits J-1 to 0 are the byte offset. Set subdivision is applied as described above to implement cache sharing, in processing block 625. In

processing block 630, cache device supports two simultaneous accesses in

cache sharing mode configured to split the cache equally. For example, CPU 111 and color/Z are allocated L/2 sets each. The number of ways allocated to each request type remains constant. In processing block 640, the physical

address for the cache portion allocated to color/Z requests are defined as with bits Z to J+K-1 are the tag address bits, bits J+K-2 to J are the set address, and bits J-1 to 0 are the byte offset. In processing block 645, color/Z requests are converted to direct-mapped cache. In processing block 650, the physical

5 address for the direct mapped cache portion allocated to color/Z requests are defined as with bits Z to J+K-1+M are the tag address bits, bits J+K-2+M to J are the set address, and bits J-1 to 0 are the byte offset. The process ends in block 699. In summary, when converting from N-way set associative to a direct-mapped cache, the set address expands by M bits. The most significant M bits
10 of the set address decode N ways. For example, to convert a 2-way set associative cache to direct mapped, M=1 for a 4-way, and M=2 for an 8-way.

| Entry Hit | New LRU Value | | | |
|-----------|---------------|----|----|----|
| | MSB | L1 | L2 | L3 |
| 0 | 1 | 1 | 1 | |
| 1 | 1 | 1 | 0 | |
| 2 | 1 | 0 | | 1 |
| 3 | 1 | 0 | | 0 |
| 4 | 0 | 1 | 1 | |
| 5 | 0 | 1 | 0 | |
| 6 | 0 | 0 | | 1 |
| 7 | 0 | 0 | | 0 |

Table 4

Table 4 shows a pseudo-LRU method to achieve an 8-way set associative cache developed from a 4-way set associative cache with 3 LRU
15 bits. The MSB set address bit is the most significant bit of the set address of the 4-way set associative cache. This technique may be applied to convert any

cache to a higher degree of set associativity. **Figure 8** illustrates the pseudo-LRU mechanism for converting to an 8-way set associative cache from a 4-way set associative cache. **Figure 8** is used to determine the entry hits as described above with reference to **Figure 5**.

5 A method and device for cache replacement in a multiple variable-way associative cache is disclosed. Although the present invention has been described with respect to specific examples and subsystems, it will be apparent to those of ordinary skill in the art that the invention is not limited to these specific examples or subsystems but extends to other embodiments as well. The
10 present invention includes all of these other embodiments as specified in the claims that follow.

CLAIMS

We claim:

- 1 1. A method, comprising:
2 partitioning a cache array dynamically based upon requests for memory
3 from an integrated device having a plurality of processors.
- 1 2. The method as claimed in claim 1, further comprising
2 subdividing one or more ways within the cache array.
- 1 3. The method as claimed in claim 1, further comprising
2 subdividing one or more sets within the cache array.
- 1 4. The method as claimed in claim 1, further comprising using
2 a single least recently used array to replace ways.
- 1 5. The method as claimed in claim 1, further comprising
2 applying a multiple pseudo least recently used update based on an entry
3 hit.
- 1 6. The method as claimed in claim 1, further comprising
2 partitioning dynamically the cache array into a direct-mapped cache.
- 1 7. A device comprising:

2 a cache memory array dynamically partitioned when multiple memory
3 requests are received from an integrated device having a plurality of
4 processors.

1 8. The device as claimed in claim 7 further comprising:
2 an integrated device having a plurality of processors connected to the
3 cache memory array.

1 9. The device as claimed in claim 7 further comprising a main
2 memory device connected to the cache memory array.

1 10. The device as claimed in claim 8 wherein the integrated
2 device includes a graphics processor and a central processing unit.

1 11. A computer-readable medium having stored thereon a
2 plurality of instructions, said plurality of instructions when executed by a
3 computer, cause said computer to perform the method of:
4 partitioning a cache array dynamically based upon requests for memory
5 from an integrated device having a plurality of processors.

1 12. The computer-readable medium of claim 11 having stored
2 thereon additional instructions, said additional instructions when executed
3 by a computer, cause said computer to further perform the method of
4 subdividing one or more ways within the cache array.

1 13. The computer-readable medium of claim 11 having stored
2 thereon additional instructions, said additional instructions when executed
3 by a computer, cause said computer to further perform the method of
4 subdividing one or more sets within the cache array.

1 14. The computer-readable medium of claim 11 having stored
2 thereon-additional instructions, said additional instructions when executed
3 by a computer, cause said computer to further perform the method of
4 using a single least recently used array to replace ways.

1 15. The computer-readable medium of claim 11 having stored
2 thereon-additional instructions, said additional instructions when executed
3 by a computer, cause said computer to further perform the method of
4 applying a multiple pseudo least recently used update based on an entry
5 hit.

1 16. The computer-readable medium of claim 11 having stored
2 thereon-additional instructions, said additional instructions when executed
3 by a computer, cause said computer to further perform the method of
4 partitioning dynamically the cache array into a direct-mapped cache.

1 17. A method, comprising:
2 converting an N-way set associative cache dynamically into a direct
3 mapped cache; including
4 removing M least significant bits from a tag address, and

5 adding the M least significant bits to M most significant bits of a set
6 address of the direct-mapped cache.

1 18. The method of claim 17, wherein N equals 2 to the power M.

1 19. A method, comprising:
2 converting an N-way set associative cache dynamically into a Z x N-way
3 set associative cache; including
4 providing Y+1 virtual copies of a pseudo-LRU array for the N-way set
5 associative cache, wherein the pseudo-LRU array is not replicated,
6 and
7 selecting a virtual copy with Y most significant bits of a set address for
8 the N-way set associative cache.

1 20. The method of claim 19, wherein Z is 2 to the power Y,
2 where Y is greater than or equal to 1.

1 21. The method of claim 19, wherein the Y most significant bits
2 of the set address for the N-way set associative cache become the Y least
3 significant bits of the tag address for the Z x N-way set associative cache.

ABSTRACT OF THE DISCLOSURE

A method and apparatus for cache replacement in a multiple variable-way associative cache is disclosed. The method according to the present techniques partitions a cache array dynamically based upon requests for
5 memory from an integrated device having a plurality of processors.





FIG. 1

| | | | | | |
|----------------------|----------------------|----------------------|----------------------|-----------------|-----------------|
| <u>309</u> | | <u>310</u> | | | |
| <u>0_A</u> | <u>1_A</u> | <u>0_B</u> | <u>1_B</u> | <u>311</u> C | <u>312</u> D |
| <u>2_A</u> | <u>3_A</u> | <u>2_B</u> | <u>3_B</u> | | |
| <u>313</u> E | <u>314</u> F | <u>315</u> G | <u>316</u> H | | |
| <u>317</u> I | <u>318</u> J | <u>319</u> K | <u>320</u> L | | |
| <u>321</u> M | <u>322</u> N | <u>323</u> O | <u>324</u> P | | |

↖ 300

FIG. 2





PSEUDO-LRU REPLACEMENT MECHANISM



PSEUDO-LRU UPDATE BASED ON ENTRY HIT

| Entry hit | New LRU value | | | | | | |
|-----------|---------------|----|----|----|----|----|----|
| | L0 | L1 | L2 | L3 | L4 | L5 | L6 |
| 0 | 1 | 1 | | 1 | | | |
| 1 | 1 | 1 | | 0 | | | |
| 2 | 1 | 0 | | | 1 | | |
| 3 | 1 | 0 | | | 0 | | |
| 4 | 0 | | 1 | | | 1 | |
| 5 | 0 | | 1 | | | 0 | |
| 6 | 0 | | 0 | | | | 1 |
| 7 | 0 | | 0 | | | | 0 |

FIG. 3






FIG. 5


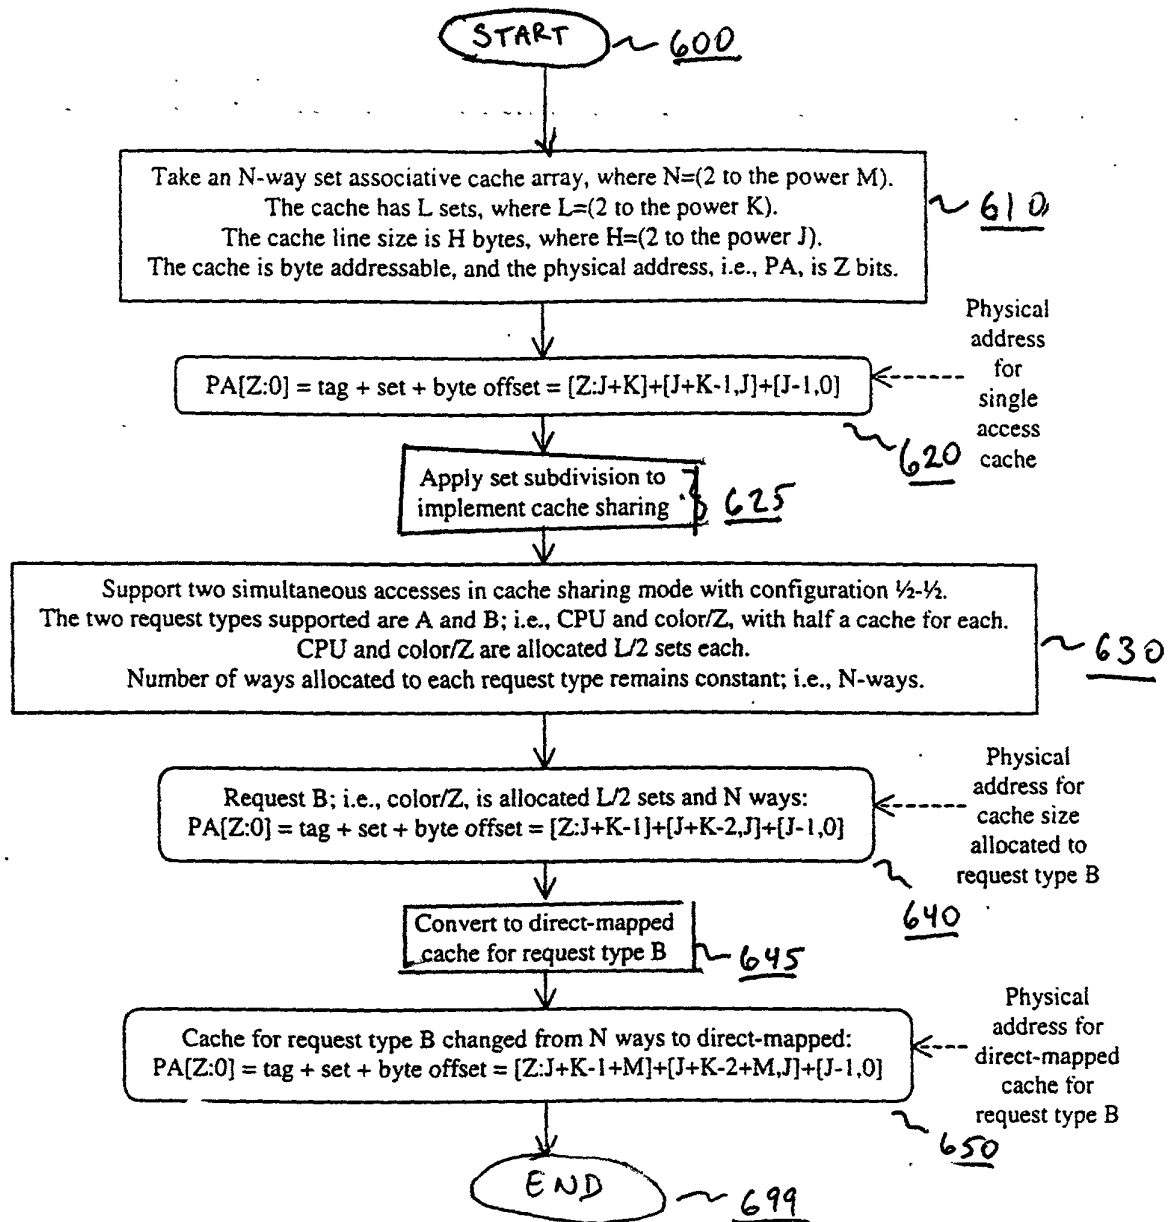



FIG. 6

\nwarrow 500

FIG. 7





PSEUDO-LRU UPDATE BASED ON ENTRY HIT

| Entry hit | New LRU value | | | |
|-----------|---------------|----|----|----|
| | MSB | L1 | L2 | L3 |
| 0 | 1 | 1 | 1 | |
| 1 | 1 | 1 | 0 | |
| 2 | 1 | 0 | | 1 |
| 3 | 1 | 0 | | 0 |
| 4 | 0 | 1 | 1 | |
| 5 | 0 | 1 | 0 | |
| 6 | 0 | 0 | | 1 |
| 7 | 0 | 0 | | 0 |

FIG. 8

-1-

| <u>Prior Foreign Application(s)</u> | | | <u>Priority Claimed</u> | |
|-------------------------------------|--------------------|---------------------------------|-------------------------|----|
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | Yes | No |
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | Yes | No |
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | Yes | No |

I hereby claim the benefit under Title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

| | |
|-----------------------------|----------------------|
| _____ Application Number | _____ Filing Date |
| _____ Application Number | _____ Filing Date |

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

| | | |
|-----------------------------|----------------------|--|
| _____ Application Number | _____ Filing Date | _____ Status -- patented, pending, abandoned |
| _____ Application Number | _____ Filing Date | _____ Status -- patented, pending, abandoned |

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Sanjeet Dutta, BLAKELY, SOKOLOFF, TAYLOR &
(Name of Attorney or Agent)
ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct
telephone calls to Sanjeet Dutta, (408) 720-8300.
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Salvador Palanca

Inventor's Signature _____ Date _____

Residence Gold River, California Citizenship Spanish
(City, State) (Country)

Post Office Address 12044 Gold Arbor Lane
Gold River, California 95760

Full Name of Second/Joint Inventor Subramaniam Maiyuran

Inventor's Signature _____ Date _____

Residence Gold River, California Citizenship Sri Lanka
(City, State) (Country)

Post Office Address 11717 New Albion Drive
Gold River California 95760

Full Name of Third/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Full Name of Fourth/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Full Name of Fifth/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Full Name of Sixth/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Full Name of Seventh/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Lisa N. Benado, Reg. No. 39,995; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; R. Alan Burnett, Reg. No. 46,149; Gregory D. Caldwell, Reg. No. 39,926; Andrew C. Chen, Reg. No. 43,544; Thomas M. Coester, Reg. No. 39,637; Donna Jo Coningsby, Reg. No. 41,684; Florin Corie, Reg. No. 46,244; Dennis M. deGuzman, Reg. No. 41,702; Stephen M. De Klerk, Reg. No. P46,503; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Sanjeet Dutta, Reg. No. P46,145; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; George Fountain, Reg. No. 37,374; Paramita Ghosh, Reg. No. 42,806; James Y. Go, Reg. No. 40,621; Libby N. Ho, Reg. No. P46,774; James A. Henry, Reg. No. 41,064; Willmore F. Holbrow III, Reg. No. P41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W. Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; William W. Kidd, Reg. No. 31,772; Sang Hui Kim, Reg. No. 40,450; Walter T. Kim, Reg. No. 42,731; Eric T. King, Reg. No. 44,188; Erica W. Kuo, Reg. No. 42,775; George Brian Leavell, Reg. No. 45,436; Kurt P. Leyendecker, Reg. No. 42,799; Gordon R. Lindeen III, Reg. No. 33,192; Jan Carol Little, Reg. No. 41,181; Joseph Lutz, Reg. No. 43,765; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Clive D. Menezes, Reg. No. 45,493; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Daniel E. Ovanezian, Reg. No. 41,236; Kenneth B. Paley, Reg. No. 38,989; Marina Portnova, Reg. No. P45,750; William F. Ryann, Reg. No. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; Joseph A. Twarowski, Reg. No. 42,191; Tom Van Zandt, Reg. No. 43,219; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Mark L. Watson, Reg. No. P46,322; Thomas C. Webster, Reg. No. P46,154; Steven D. Yates, Reg. No. 42,242; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Firasat Ali, Reg. No. 45,715; and Justin M. Dillon, Reg. No. 42,486; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and Alan K. Aldous, Reg. No. 31,905; Edward R. Brake, Reg. No. 37,784; Ben Burge, Reg. No. 42,372; Jeffrey S. Draeger, Reg. No. 41,000; Cynthia Thomas Faatz, Reg. No. 39,973; John N. Greaves, Reg. No. 40,362; Seth Z. Kalson, Reg. No. 40,670; David J. Kaplan, Reg. No. 41,105; Peter Lam, Reg. No. 44,855; Charles A. Mirho, Reg. No. 41,199; Leo V. Novakoski, Reg. No. 37,198; Thomas C. Reynolds, Reg. No. 32,488; Kenneth M. Seddon, Reg. No. 43,105; Mark Seeley, Reg. No. 32,299; Steven P. Skabrat, Reg. No. 36,279; Howard A. Skaist, Reg. No. 36,008; Gene I. Su, Reg. No. 45,140; Calvin E. Wells, Reg. No. P43,256; Raymond J. Werner, Reg. No. 34,752; Robert G. Winkle, Reg. No. 37,474; and Charles K. Young, Reg. No. 39,435; my patent attorneys, of INTEL CORPORATION; and James R. Thein, Reg. No. 31,710, my patent attorney with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

APPENDIX B

Title 37, Code of Federal Regulations, Section 1.56 Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) Prior art cited in search reports of a foreign patent office in a counterpart application, and
 - (2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.
- (b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and
- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
 - (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
 - (2) Each attorney or agent who prepares or prosecutes the application; and
 - (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.
- (d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.